

an evaluator configured to determine a performance score for the first model and a performance score for the second model.

2. The apparatus of claim 1, wherein:
the orchestrator is further configured to deploy the second model to the decision engine, the first model being deployed to a first set of one or more users and the second model being deployed to a second set of one or more users; and
the decision engine is further configured to provide the optimized action to the decision agent.

3. The apparatus of claim 1, wherein:
the decision engine is further configured to service one or more action requests from the decision agent during the test period; and
an archiver configured to store transaction data, wherein transaction data comprises at least one or more action requests, one or more optimized actions, or responses to the one or more optimized actions.

4. The apparatus of claim 1, wherein:
the evaluator is further configured to determine transaction data of the first model and the second model from the test period; and
the evaluator is further configured to determine the first model performance score and the second model performance score using a statistical hypothesis test.

5. The apparatus of claim 1, wherein:
the orchestrator is further configured to allocate traffic to the first and the second model during the testing period; and
the orchestrator is further configured to alter the allocation of traffic to the first and the second model.

6. The apparatus of claim 1, wherein:
the evaluator is further configured to simulate an external system using stored transaction data, the simulated external system comprising the decision agent and an environment; and
the evaluator is further configured to determine an average cumulative net response for the first model and an average cumulative net response for the second model.

7. The apparatus of claim 1, wherein the orchestrator is further configured to store one or more parameters, wherein a parameter comprises at least one of a parameter to indicate a relative split of traffic across the models used during the test period, a parameter to specify a confidence level when comparing a performance of the first model and the one or more secondary models, and a parameter corresponding to a desired threshold to compare an average performance of the first model and the one or more secondary models.

8. The apparatus of claim 1, wherein the model builder is further configured to build the first model using transaction data collected from a legacy system or generated from a rule-based system.

9. A continual learning method, implemented by one or more processors, the continual learning method comprising:
receiving one or more action requests from a decision agent;
deploying a first model to a decision engine;
initiating an observation period;
building a second model, wherein the second model comprises collected transaction data from the observation period;
initiating a test period;

determining a performance score for the first model and a performance score for the second model; and
selecting a model providing an optimized action.

10. The method of claim 9, further comprising deploying the second model to the decision engine, wherein the first model is deployed to a first set of one or more users and the second model is deployed to a second set of one or more users; and

providing the optimized action to the decision agent.

11. The method of claim 9, further comprising:
servicing one or more action requests from the decision agent during the test period; and
storing transaction data, wherein transaction data comprises at least one or more action requests, one or more optimized actions, or responses to the one or more optimized actions.

12. The method of claim 9, wherein determining the performance score for the first model and a performance score for the second model comprises:

determining transaction data of the first model and the second model from the test period;
determining the first model performance score and the second model performance score using a statistical hypothesis test.

13. The method of claim 9, wherein determining the performance score for the first model and the performance score for the second model comprises:

allocating traffic to the first and the second models during the testing period; and
altering the allocation of traffic to the first and the second models.

14. The method of claim 9, wherein determining the performance score for the first model and the performance score for the second model, further comprises:

simulating an external system using stored transaction data, wherein the simulated external system comprises the decision agent and an environment; and
determining an average cumulative net response for the first model and an average cumulative net response for the second model.

15. A non-transitory computer readable storage medium, implemented by one or more processors, storing a continual learning system for causing a computer to function as:

a decision engine configured to receive one or more action requests from a decision agent, and to select a model providing an optimized action;

an orchestrator configured to deploy a first model to a decision engine, to initiate an observation period, and to initiate a test period;

a model builder configured to build a second model, wherein the second model comprises collected transaction data from the observation period; and

an evaluator configured to determine a performance score for the first model and a performance score for the second model.

16. The non-transitory computer readable storage medium of claim 15, wherein:

the orchestrator is further configured to deploy the second model to the decision engine, the first model being deployed to a first set of one or more users and the second model being deployed to a second set of one or more users; and

the decision engine is further configured to provide the optimized action to the decision agent.